



Hybrid Cloud

**Cloud-native meets
hybrid cloud:
A strategy guide**

Building application environments for
reliability, productivity, and change

03

Introduction

04

Part one: Modern IT challenges

- Transforming Red Hat IT
- Fast or stable?
- Creating and evolving an effective hybrid cloud application environment

07

Part two: Strategy primer

- Situation assessment and objectives
- Considerations for policies and guide rails
 - Decentralized authority
 - Universal capabilities
 - (Some) constraints = freedom
 - Self-service provisioning
 - Products over projects
 - Tend toward antifragility
 - Automation as code
 - Security in depth
 - Bringing policy areas together
- From policy to actions—some advice

21

Part three: Architecture for success

- The big picture
- Platform and delivery
- Applications from custom code
- Integration is critical
- Applications from process automation

32

Part four: The application environment journey

- Fragmented bureaucracy
- Steps along the cloud-native and hybrid cloud path
- Is there a methodology?

35

Part five: Conclusions and next steps

Introduction

Information technology (IT) systems are the glue that holds modern enterprises together. They power internal systems, are often critical market differentiators, and need to deliver ever-improving functionality for customers, employees, and partners. At the same time, system complexity and speed of change have escalated: deployments routinely span multiple on-premise and cloud locations, an array of software technologies are used, and the teams behind these systems can vary greatly in their skills. The increased focus on digital transformation highlights the business advantages of change, but often does not address how to deliver real value with new technologies.

These pressures lead to a fundamental tension between the reliability and productivity of the IT environment. On one hand, systems must be kept functioning, secure, and predictable in their operation. At the same time, developers and operations teams need to be able to evolve the systems rapidly to deliver new or improved functionality to users so the whole organization benefits from progress.

This core tension between reliability and productivity impacts many choices, but it becomes particularly problematic given current pressure for rapid change and improvement. IT teams find themselves continually having to balance short- and long-term change in application delivery and development. For example, trading the use of existing systems and infrastructure to develop new functionality versus redesigning development and delivery capabilities. The long-term aim is to keep adding value with the company-wide application environment, while improving agile development and delivery.

Managing these pressures is relevant since IT systems have become highly distributed across datacenter locations and public clouds—hybrid cloud IT is a reality for most organizations. In addition, a myriad of new development technologies in cloud-native and other areas open up huge new opportunities to make development teams more productive.

This guide is designed as a strategy primer for enterprise architects and IT leaders to map out how to tackle modern IT strategy. It is focused on the convergence of three trends:

1. The importance of IT systems to modern organizational success.
2. Hybrid cloud and the shift to multiple datacenters, platforms, and cloud locations.
3. Cloud-native application development technologies and how they can be combined with existing approaches to create a robust, productive IT system.

This e-book covers context considerations, strategy challenges, and high-level architectural considerations with case study examples.

This e-book:

- ▶ **Focuses on strategic aspects and high-level architectures**, specifically, high-level concepts that can be applied in a wide variety of contexts. It does not address detailed deployment considerations. Resources for deeper technical analyses are offered throughout.
- ▶ **Does not assume greenfield development only.** While some teams may have the luxury of building new systems from scratch, the majority of our advice applies to the evolution of complex existing systems.

"Like many of our customers, we faced the real questions of, 'How can we do better than just keep up?' and 'What can we do to be well prepared for whatever comes our way?'"

Mike Kelly
CIO, Red Hat

- ▶ **Provides general technical applicability.** Red Hat® customer and technology examples are used in a variety of areas. While we believe Red Hat technologies are right for the job, most approaches discussed in this e-book can be applied in other technology environments or using a mix of Red Hat and other technologies.

Companion content can be found in the Red Hat [cloud-native solutions](#) and [topics pages](#).

Part one: Modern IT challenges

When he took the role of Chief Information Officer at Red Hat, Mike Kelly was in a difficult position. Red Hat had a robust IT infrastructure that had grown tremendously, but even more growth was approaching. Red Hat's rapid growth meant that the IT team needed to get ahead of the growth curve while still operating within a constrained budget.

How did Mike and his team approach these challenges?

Transforming Red Hat IT

IT systems can be thought of as the central nervous system of an organization, carrying information, coordinating actions, and executing many of the tasks that allow the organization to function. The size, density, and complexity of this system of applications has increased for most organizations in recent years. Red Hat's experience has been no different.

In 2016, Red Hat operated nearly 1,000 unique, independent applications and services that served various parts of the business. That number was too high. These applications were run by different teams, on different technology stacks, and across a number of redundant datacenter locations. Additionally, the applications powering Red Hat's day-to-day business ranged from custom Java™ and other applications, to integration of Software-as-a-Service (SaaS) and data stores, to complex workflow processes.

These challenges required the team to focus on three key objectives:

1. Match the speed and adaptability demands of digital business.
2. Improve the availability, resiliency, and security of digital systems.
3. Continue to reduce operating costs.

In combination, these objectives pose a typical dilemma: How to fulfill the need for speed, adaptability, availability, and security—but with lower operating expenses.

The project begins

[Red Hat Open Innovation Labs](#) is an immersive program that uses residency-style engagements that pair engineers with open source experts, transforming customer ideas into business outcomes.

To get started, Red Hat IT engaged the [Red Hat Open Innovation Labs](#) team to receive the same thorough analysis routinely done for Red Hat customers. The Open Innovation Labs program provides immersive residencies that help organizations solve real business challenges with a combination of technology and process changes. The result for Red Hat IT was comprehensive insight into the choices they faced, options to address some of the most pressing issues, and a roadmap for continued process refinement over time.

This analysis spurred on change across a broad range of areas in Red Hat's IT processes, including:

- ▶ A shift from a multiple-location, on-premise failover infrastructure to a hybrid cloud infrastructure, including first one, then two public cloud zones from different providers.
- ▶ A move to a container-based, cloud-native deployment and application development mode, with containers as the primary unit of deployment.
- ▶ Significant shifts in software development technologies and methodologies to allow a wide variety of programming languages to be used according to the needs of each application.
- ▶ Significant cost savings from containers and the ability to move from a parent-child failover setup to active-active-active cluster setups for greater reliability and faster updates.
- ▶ Upgrades to a range of applications and the increased adoption of application programming interfaces (APIs) and agile integration as a way to connect these applications.

Observable, measurable outcomes

By September 2021, the transformation had exceeded early goals but was still seen as an ongoing process that continues to evolve as new challenges and opportunities arise. In addition to achieving a more than 55% smaller footprint per application where measurable, containerized applications rose to represent 35% of the portfolio, and technical debt was reduced by retiring 29% of the on-premise resources. The team moved from 150 racks of equipment to approximately 20.¹ Overall, the use of bare metal for application hosting was significantly reduced.

There are now numerous ways for IT team members—and anyone in the company—to deploy new applications and share access with others. Infrastructure for production systems is resilient and lower in cost.

The team also learned to navigate the internal cultural challenges inherent to change. It turned out to be just as important to value inputs from skeptics on the team as those from the enthusiastic technology adopters.

Read more about [Red Hat IT's transformation](#).

Fast or stable?

As the Red Hat IT story shows, new technologies can provide genuine benefits for addressing modern IT challenges, but their application takes thought and coordination. The question of how to navigate short- and long-term change is particularly acute in large organizations since development teams need to move fast and adopt new and more productive technologies, while operations teams are responsible for reliability and stability.

¹ Interview with Seth Shevach. November, 2021.

The DevOps movement has made this conflict visible to different groups within the organization. However, even with DevOps adoption, there are typically varying practices in every organization. The challenge is exacerbated by a wide range of high-pressure trends that diversify technology choices:

- ▶ The vast majority of enterprises (92% according to a recent Flexera/Rightscale report²) are pursuing a multicloud strategy. In many cases, organizations have become hybrid cloud by default, simply because different groups adopt different cloud providers.
- ▶ Programming languages continue to explode in use, both in general and in the enterprise. While Java and Javascript still dominate in the enterprise, other languages are gaining traction in specialized applications. (For programming language rankings, see Red Monk's June 2021 public data survey.³)
- ▶ The move toward containerized infrastructure and OpenStack® means that there are fewer organizations using bare-metal servers.⁴

The challenge is how to balance these pressures.

Creating and evolving an effective hybrid cloud application environment

We use the term application environment to refer to an organization's set of capabilities for application delivery and development. This broad view is deliberate, and it underlies the holistic nature of the modern IT challenge. Specifically, legacy and new applications need to work together, new technology needs to add value without disrupting reliability, and IT systems are found in an ever-increasing variety of locations.

An organization's application environment undergoes constant change as new applications are deployed, tools and processes are updated, and applications generate data and transactions. Ultimately, the production applications that the environment supports are what add value to organizations employees, partners, and users.




Figure 1. Application environments support applications that deliver value

2 Flexera, "[RightScale 2021 State of the Cloud Report from Flexera](#)," 2021.

3 O'Grady, Stephen. "[The RedMonk Programming Language Rankings: June 2021](#)." Red Monk blog, November 16, 2021, .

4 Tripathi, Vivek Mani. "[Digital transformation: 7 in-demand technology skills](#)." The Enterprisers Project, November 10, 2021.

Thinking of an organization's application environment as a single, organization-wide, evolving ecosystem is a powerful metaphor since it brings into focus the properties the environment needs to exhibit and what tensions are in play as individual decisions are made.

Many of the compromises mentioned in later sections, including reliability versus productivity, predictability versus the ability to change, security versus convenience, performance versus cost, and rigidity versus freedom, look like an either-or proposition. They look like choices to be made in opposition to one another. While trade-offs do need to be made, for an organization to excel, IT still needs to deliver on both aspects of each compromise:

- ▶ Reliability and productivity
- ▶ Predictability and ability to change
- ▶ High performance and low cost

With deliberate, strategic planning and focused decisions, it is possible to deliver on many of these needs.

Part two: Strategy primer

The process of developing and adopting a strategy is ever-evolving and can pull from several different information sources and philosophies. At its essence, the process requires breaking strategy development down into three steps. These three steps provide a broad overview of what types of decisions comprise a long-term strategy.

1. **Assess the situation and set objectives.** Capture what's happening now—the challenges and opportunities presenting themselves, as well as significant constraints on action. It would be typical to include both a long-term "where do we want to be in 5-10 years?" element and a short-term "what's coming in the next 12-18 months?" view.
2. **Set policies and guidelines.** Capture broad principles and rules of engagement that determine the overall direction of the strategy, what is in and out of scope and general practices for how certain things are done. Policies and guide rails are typically guidelines that will remain relevant and useful in the midterm: at least a single year, but preferably four-five years.
3. **Determine what to do next.** Create a set of specific actions to take right now that bring the organization closer to key objectives and more in line with the established policies and guidelines.

In this section, we cover the first two areas of this structure (situation assessments and considerations for setting policies). The third area is covered in depth in [Part three](#) (architecture) and [Part four](#) (where to start).

As you will see in later sections, one of the recurring tactics for success is the use of decentralized authority. As a result, it is likely that any strategy will be most effective as a multilevel construct with broad policies and guide rails at the organizational level, with leeway for individual groups to add their own local instantiations to suit their own needs.

Situation assessment and objectives

In conducting a situation assessment of your organization's IT environment, there are several factors to consider, including existing policies, personnel, and projected usage. Which aspects are most critical to address will depend on and vary by each organization, but at a high level of abstraction, there are a few elements of the overall application environment that are often among the most important to consider.




Figure 2. Simplified IT landscape view

Specifically, every organization is likely to have:

- ▶ **A set of on-premise datacenters—infrastructure hosted with one or more private hosting services, and public cloud providers.** The platforms are often distributed across geographies for failover or data protection. Regions often comprise one or more types of operating systems, virtualization systems, private cloud deployments, or containerized infrastructure.
- ▶ **A set of applications and application development solutions.** From an application perspective, the set is likely to include custom code written by the organization or third-party consultants, commercial off-the-shelf (COTS) solutions, SaaS solutions, and a range of other applications. In addition, the organization likely has application servers or other technology to host custom code, integration solutions, messaging, and process management to connect individual applications together.
- ▶ **Developer tools, processes, automation, and management capabilities** that manage different parts of the infrastructure.
- ▶ **A range of stakeholders** who depend on the application environment to be productive, including developers and operations teams who create and operate applications, as well as employees, customers, and partners who use the resulting application.

Sketching out these major components is a useful exercise since it provides a view of the breadth and depth of the whole application environment.

In terms of objectives, we will focus on two near-universal areas: reliability and productivity.

However, there may be other high-priority areas specific to your organization that you need to consider. We expect some of the guidance that follows will address other common areas as well, but we focus on reliability and productivity since they are common across organizations.

What we mean by “productivity”

For our purposes, productivity refers to the entire organization and all its employees. How much is getting done throughout the whole organization? In practice, we focus primarily on the productivity of the IT, development, and operations teams responsible for improving the applications available to their end users. However, by extension, these application improvements are a large part of what delivers productivity in the remainder of the organization. There are areas, as we will see in [Part three](#), where the definitions between developers and users are becoming blurred.

Set clear objectives

We encourage thinking through three types of objectives:

- 1.** General outcomes. What do we want the application environment to do better?
- 2.** Common patterns and antipatterns. What are examples of successful applications or teams that we would like to replicate? What are examples of failures to avoid?
- 3.** Gains in specific projects or initiatives. What specific activities, such as new features, requests for new technology, and retiring old functionality, are underway but might benefit from a shift in direction?

Considerations for policies and guide rails

Once objectives are set, the challenge becomes creating processes that will lead to achieving goals. Generally, this works best by establishing a set of well-informed, clear policies and guide rails that different groups in the organization can agree upon and follow.

The precise set of policies will be different for each organization. In this section, we cover a number of policy area considerations that we have seen work in large IT environments. The set covered here is a distillation of our experience working with a wide range of customers worldwide and an adaptation of prevailing wisdom from the field that has proven useful in practice.

The policy areas covered are applicable at the organizational level. Although you may have seen some of these strategies in other hybrid cloud guides, they are more abstract and do not necessarily relate directly to the adoption of specific technologies. Later sections will provide technology-focused advice on implementation.

Decentralized authority

The trend toward decentralization in system design has been prevalent for a number of years, and in technology, it is most strongly expressed in methodologies such as microservices. This architecture makes sense for many types of systems, and it is one of the areas where cloud-native technology has provided major breakthroughs.

Before considering implementation, there is another level at which decentralization can play a key role.

Conway's Law states that "organizations which design systems (in the broad sense used here) are constrained to produce designs which are copies of the communication structures of these organizations."⁵

Melvin Conway

Decentralized systems are not appropriate for every part of an IT system, but decentralization in organizational terms is almost always beneficial. It is more important to structure responsibility boundaries among teams than code boundaries. This approach may seem like a straightforward invocation of Conway's Law; however, it is often done incorrectly in organizations as they rush to implement decentralized application development models, such as microservices.

We express this policy principle as: Centralized authority -> Decentralized authority.

Enablement rather than ownership

Centralized IT can empower other teams throughout the organization and share the responsibility for development and operations, which works better than trying to centralize authority.

This organizational approach does not mean that central teams devolve all responsibility. There are key processes and systems they must require other teams to follow. However, by providing the means for others to take responsibility, groups are much more engaged in making the whole organization function, which leads to:

- ▶ Reliability gains and better sharing of best practices (smaller teams may uncover problems before central teams).
- ▶ More individual investment in operational outcomes.
- ▶ Clearer supervisory focus on what is important versus what is not.

Productivity gains include:

- ▶ Better ability of local teams to rapidly adapt regulations to their context.
- ▶ Faster decision-making for smaller items, resulting from a greater sense of autonomy and ownership.
- ▶ More pride in execution and sense of purpose, leading to higher levels of dedication and even better productivity.

Decentralization may be supported by technology via remote work, efficient ways of tracking decisions, and organizing teams, but it is first and foremost a cultural issue. Small shifts in directives and responsibilities can lead to significant differences in outcomes.

Universal capabilities

Increased numbers of public cloud deployments and datacenters can result in a rapid proliferation of new segmentation in an organization:

- ▶ Certain services only available on Microsoft Azure, others on Amazon Web Services (AWS), and others only in the corporate datacenter
- ▶ Certain services only available in specific areas of a hosted system controlled by a specific group

⁵ Conway, Melvin. "[How Do Committees Invent?](#)" Datamation, April 1968.

Some of this separation exists for good reason—such as security firewalls or data protection—but determining where to host a new piece of functionality often involves an extremely complex process. Other types of dysfunction, including difficult learning curves on how to do the same task in different environments, are not far behind.

The key policy shift to consider is a move from the typical minimalist approach to solving a challenge within a single group to addressing similar instances across an organization, going from:

"We needed capability X, we selected or built it in environment A, and if you need to use it, it's there."

And shifting to: "We needed capability X, we selected or built it and made it available in environment A. If required, it could be rolled out with the exact same interface or configuration in environments B, C, and D."

In other words: Single-use functionality -> Universal capabilities.

This concept seems simple and obvious, yet the forethought required is only done briefly, if at all.

As a general principle, capabilities in a hybrid cloud environment should:

- ▶ Be available in an identical form in as many datacenters and cloud locations as may be needed. They should use the same APIs, endpoints, and configurations whenever possible.
- ▶ Have distributed capabilities such as messaging, integration, logging and tracing, and process automation that run efficiently across datacenter and cloud locations, and transparently connect those locations.

Note: Depending on the service, being "available" in another datacenter or cloud may mean that either the same instance is simply accessible as a service of an API remotely from other locations or that local instances of the same service are available.

By insisting on interface uniformity and the ability to recreate capabilities in the exact same way across locations, it is much easier to create failover configurations. Doing this reliability prework as new capabilities are planned and configured has a dramatic effect later in the process.

This approach also leads to a uniformity of execution and reuse that delivers greater benefit than using different underlying capabilities in different locations. From a developer productivity point of view, having the same set of available services reduces the amount of learning required and increases the applicability of skills learned by the team.

(Some) constraints = freedom

The desire to empower a greater number of teams in the organization to contribute and operate autonomously is fine in principle. But how does it affect the system in practice? The most important points become evident when users raise concerns in terms of control versus freedom: "You tell us we're responsible and autonomous, yet you impose a huge raft of restrictions and technologies." Each end of the spectrum poses significant challenges:

- ▶ Too little enforcement can lead to a lack of rules and experiments that create fragile application environments.
- ▶ Too much enforcement tends to lead to resentment, disengagement, or creative workarounds that end up being even more fragile. These workarounds often contain useful ideas but end up impeding long-term progress.

To succeed, the question is not "how restrictive should our policies and guide rails be," but "where should we be restrictive and where should we be flexible?"

A useful illustration of this approach is driving regulations. For example:

- ▶ The side of the road to drive on is typically strictly mandated.
- ▶ The type of vehicle that can be used, its color, power, or model, can be chosen almost arbitrarily by the driver.

So the strict application of one rule (the side of the road to drive on), while technically a restriction on freedom, actually empowers drivers since it lets them know what side of the road everybody else will be driving on. In practice, there are some restrictions on the vehicle type as well, but they are generally broad enough that drivers feel free. So in reality, many of the strict driving rules free up capacity to choose and the ability to innovate in other areas.

In IT systems, much of the challenge is establishing where to be strict and where to allow freedom. A way to express this as a policy dimension is: Governance barriers -> Agreed interfaces.

Working to establish which interfaces and core processes should be strict and adhered to by everybody benefits reliability by:

- ▶ Enforcing modeling discipline in key areas.
- ▶ Requiring communication about these interfaces and technology choices between groups.
- ▶ Creating a set of fixed points to measure drift, change, and degradation.

Surprisingly, this approach can also lead to increased productivity:

- ▶ Interfaces at boundaries make more radical change within domains possible (e.g., the adoption of a variety of implementation technologies).
- ▶ Teams waste less time coordinating solutions to preventable issues with other groups when potential problems are fixed ahead of time.

"The partnership between Red Hat and Amadeus is a win-win collaboration: Amadeus has contributed to the improvement of Drools' internal algorithms and test coverage, and Red Hat has helped us to build a rule engine that better meets our needs."

Matteo Casalino

Technical Expert Business Automation, Rule and Inference Engines, Amadeus

6

Self-service provisioning

Setting the goal of enabling IT capabilities everywhere, independent of the cloud location in which they might initially be needed (policy area two), can be thought of as a horizontal dimension to creating a uniform, consistent environment. However, there is also a vertical dimension to this notion—a view on how a capability is made available. A simple way to think about this approach is as a series of questions:




Figure 3. How self-service oriented are your IT systems?

This spectrum for provisioning addresses two things:

1. The amount of time and effort required by team members to get set up with the resources they need.
2. The potential skill levels that team members require to access certain resources.

We express this continuum as: Applications → Self-service services.

Many of the capabilities available in IT environments are installable applications that are different from test to production, are only available in certain places, and may require significant skills to set up correctly. The more self-service oriented the applications and the application development tools are, the more productive the organization is likely to become. For example, organizations with more self-service capabilities:

- ▶ Reduce the time required for developer teams to provision resources, evaluate options, and become productive.
- ▶ Allow individuals to contribute faster by removing the initial barrier of laborious system setup that often requires specialist knowledge.

There are also clear reliability benefits that come from using this approach:

- ▶ Fewer errors and debugging due to the automation of service delivery—especially during initial setup and configuration where errors can lead to undetected problems later, breaking up production execution
- ▶ Higher service availability when provisioning for many users across a larger, shared resource base than for an individual user
- ▶ Better capacity and budget planning that leads to more reliable and efficient resource use

6 [Red Hat Amadeus presentation](#) at 2020 Red Hat Summit, April 2020.

Products over projects

When adapting IT operation processes, the scope of change should not be limited to a single, or even multiple, finite projects. Adaptation and change must be baked into the resulting product itself. In the high-pressure, continual operation of IT systems, completing a defined list of project tasks or milestones can be motivating. This approach makes sense for one-off issues or tasks, but it is not ideal when the deliverable is meant to be a durable system element on which other teams rely.

One way to think of this concept is as the difference between hunters and farmers. As a hunter, it feels rewarding to capture the game needed for that day, much like reaching a project's completion. In contrast, farmers need a regular operation that produces results over time, for this season and the next. Thinking about the whole application environment, real progress is the improvement of core assets, or the creation of new ones, which meaningfully upgrade and improve production for the long term.

We express this policy principle as: Projects -> Products.

Thinking of IT systems as products causes us to consider ongoing production, the users, and what the users need. Then, if you are planning to make changes, how to best communicate them ahead of time. These are all considerations that would be applied to products used by external customers. But this thinking is equally valuable—even critical—to apply to any system that affects any user group, including:

- ▶ Partners.
- ▶ Employees.
- ▶ Developers.
- ▶ Operations teams.

Without product thinking, system reuse is not likely to happen. But a commitment to product thinking can:

- ▶ Increase reliability by creating clearer ownership of reusable services and capabilities, create more predictable upgrade-downgrade-retirement cycles, and result in fewer surprises for operations and other teams.
- ▶ Increase productivity by strengthening the services and capabilities on which teams rely, reduce the need to rebuild or replicate locally, and allow internal services to potentially become valuable and tested external offerings in the future.

Tend toward antifragility

The “design for failure” principle has been popularized with cloud-native development. These approaches are arguably an IT overview of a more general notion of “antifragility,” coined by the writer Nicolas Nassim Taleb. In a high-level explanation, Taleb describes the concept: “Antifragility is beyond resilience or robustness. The resilient resists shocks and stays the same; the antifragile gets better.”⁷

To put it more simply, fragile systems would be ones built like wine glasses. A broken wineglass is impossible to get back in working order. An antifragile system would be more like the human body. A certain amount of stress, such as through exercise, can make it stronger, not weaker, in the end.

⁷ Taleb, Nassim Nicholas. *Antifragile: Things That Gain from Disorder*. United States: Random House, 2012.

The more subtle, yet important, point behind antifragility is that human instinct—and the prevalent instinct in IT—is to focus on the protection of fragile systems over programs for continual improvement. We generally aim to prevent problems in the first place, rather than looking at how we would recover and grow from them when they inevitably happen. This tendency is one of the reasons these approaches are still rare in practice. Most IT systems are arrays of isolated, brittle systems with protections in place to prevent random elements from causing a disturbance.

Seen from an IT perspective, antifragility really comprises two distinct functions:

1. Observability—the ability to detect failures or certain bad behaviors.
2. Healing—the ability to react appropriately to reduce damage and come back stronger.

Either or both of these stages may involve human intervention, although automation is becoming increasingly critical to react quickly.

The key mindset shift is to one that considers the range of failure modes that we can currently tolerate. And then, how do we detect them? How do we analyze data from failures? What systems should we put in place to grow as failures occur? How might we automate the process?

We express this policy principle as: Protected systems -> Antifragility.

This axis is probably the most challenging for an IT organization to move along, and it is arguably the most valuable. There are huge gains to be made, even with gradual progress toward this goal. From a reliability perspective, looking at mechanisms to systematically improve beyond each failure gradually makes the system more robust. From a productivity perspective, less worry about breakage allows for faster experimentation cycles.

A final, important point is that while fragility and antifragility might be initially thought of in terms of systems and operations, these concepts are just as important for development processes, product development, and team collaboration. Any process that contributes to forward motion can benefit from some progress along the antifragility axis.

Automation as code

Automation has been noted as an important component in two previous policy areas, but it is so critical that we also wanted to focus on it explicitly.

Most IT systems today would not function if there was not already automation in place. But in many cases, automation designed in previous periods of technical change can become part of the problem—for example:

- ▶ Bespoke apps with only one function, written in a programming language that is no longer common in the organization—or the world.
- ▶ Complex, interwoven conjunctions of scripts that no one is willing to touch lest they fail.

What is the difference between good automation and bad automation? While Red Hat's [The automated enterprise e-book](#) describes automation best practices in detail, there is one general principle to understand: the difference between implicit and explicit automation.

Implicit automation is one-off code or scripting executing in place, as opposed to explicit automation that is recognized as code and is versioned, tested, updated, and managed explicitly as part of system configuration.

We express this policy transformation as: Handcrafted configuration \rightarrow Automation as code.

Almost every organization has handcrafted configurations in many parts of its systems. Over time, transforming these to be explicitly managed code creates a lot of value. For example:

- ▶ From a reliability perspective, it reduces the chance of error due to missed dependencies, speeds up changeability, and creates discipline in change processes.
- ▶ From a productivity perspective, it reduces the areas of systems that are poorly understood and lessens the amount of time spent on hard-to-diagnose failures.
- ▶ Automation as code also increases reusability and lowers the learning curve.

Overall, moving to automation as code is a prerequisite for true DevOps success because it separates configuration from code.

Security in depth

Perimeter security has been an insufficient approach to IT security for a long time, yet many organizations still rely on it, either for the whole infrastructure or within subgroups of datacenters. In hybrid cloud scenarios, not only are applications distributed between multiple datacenters and clouds, there are also integrations and channels for data flow between some of these locations. This configuration means that compromised systems in one location could lead to unwanted access in others. Therefore, it is worth assessing security in at least three dimensions:

- 1. Vertical.** From the operating system up through virtualizations, container management, and code execution runtimes to the management of applications, good security requires consistent and up-to-date software in the entire stack.
- 2. Horizontal.** With microservices patterns, APIs, and communication between applications, good security requires encryption, tracking, and access control for most, if not all, communication.
- 3. Team and life cycle.** Even if the right technologies and procedures have been researched and chosen, development and operations teams need to be able to apply those technologies and processes. As much as possible, these processes should be automated and enforced as part of the development life cycle.

One shift to consider is whether to maintain separate security and development teams. Development teams nominally stick to security guidelines, but are typically only required to consider security issues in the short period before production code deploys. Alternatively, a collaborative, more distributed approach may be better, allowing:

- ▶ Development, operations, and security teams to collaborate much more closely using a DevSecOps approach.
- ▶ Security protections to be built in from the start at the prototyping stage.
- ▶ As many core security procedures as possible to be automated and hard-coded into development and operations environments, making security decisions easier.

We express this policy principle as: Perimeter security -> Pervasive security.

Clearly, security has an important impact on the reliability axis of IT systems. But thinking through security models, expanding the responsibility across teams, and automating many of the processes (like preconfigured, permitted container images), means more preemptive planning takes place—for an even greater reduction of risk. Protecting processes and data integrity, pervasive security reduces downtime and the risk of catastrophic loss.

Security can be seen as a drag on productivity because it often involves multiple, complex development steps rather than an embedded password, direct database access, or some other shortcut. While this may be true in the short term, with enough automation and support for processes, security can improve productivity through reductions in future refactoring and time lost responding to incidents. Having well-established and reusable security mechanisms ultimately makes developers' lives much easier. Therefore, reducing the complexity in setup and configuration is time well spent.

Bringing policy areas together

Not all policy areas will be of equal importance to every organization, and in some areas, different choices might make sense. However, policies broadly cover key areas where significant benefits can be unlocked when looking holistically at IT systems—see Table 1.

Table 1. Overview recommendations with side-by-side benefits

Recommendation	Reliability gains	Productivity gains
Decentralized authority	<ul style="list-style-type: none"> ▶ Best practices are easier to apply ▶ Control and supervision of key items ▶ Focus on the most important issues 	<ul style="list-style-type: none"> ▶ Flexibility to make the best of local context and knowledge ▶ A greater sense of autonomy and ownership
Universal capabilities	<ul style="list-style-type: none"> ▶ Serve from multiple locations—built-in failover ▶ Uniformity of execution ▶ Speed of services 	<ul style="list-style-type: none"> ▶ Convention over configuration—tools and processes work the same way everywhere ▶ Transferable skills
Constraints = freedom	<ul style="list-style-type: none"> ▶ Modeling discipline ▶ Explicit communication about interfaces ▶ Fixed reference points to measure drift, change, and degradation 	<ul style="list-style-type: none"> ▶ Fixed interfaces at boundaries support more radical change within domains (e.g., polyglot programming) ▶ Less time wasted across groups when interaction patterns are agreed upon in advance
Self-service	<ul style="list-style-type: none"> ▶ Automation of service delivery and scaling up or down to avoid errors ▶ Greater service availability ▶ Better capacity planning and efficiency 	<ul style="list-style-type: none"> ▶ Speed of setup ▶ Allowing users with less technical specialty knowledge to use some services
Product mindset	<ul style="list-style-type: none"> ▶ Greater and clearer ownership of services ▶ Predictable cycles for upgrades, downgrades, retirement, etc. 	<ul style="list-style-type: none"> ▶ Stronger components to reuse ▶ No need to rebuild locally ▶ Internal and external use of the same product is more efficient
Tend toward antifragility	<ul style="list-style-type: none"> ▶ Big improvements in reliability—learn from every failure and fix ▶ Grow tolerance of breakage range ▶ Decouple dependencies to avoid cascading failures 	<ul style="list-style-type: none"> ▶ Less worry about breakage ▶ Faster experimentation cycles
Automation as code	<ul style="list-style-type: none"> ▶ Reduce error by removing human steps required ▶ Speed up change when fixes need to be deployed ▶ Create discipline in designing change processes 	<ul style="list-style-type: none"> ▶ Remove or speed up repetitive tasks ▶ Less failure means more time for innovation
Security in depth	<ul style="list-style-type: none"> ▶ Fewer security incidents improve uptime ▶ Preemptive planning for issues at multiple levels gets more team members thinking about security 	<ul style="list-style-type: none"> ▶ Include security into some systems from the outset ▶ Shared responsibility may slow things down initially but eventually speeds up the whole

From policy to actions—some advice

Once policies and guide rails are established, the next step is to determine actions that move the organization and deployed systems forward to the desired state. These actions can be directed toward the general outcomes planned, reinforcement of patterns, or specific projects. The next two sections offer considerations for this critical implementation phase, but we will start with some general advice:

- ▶ **Focus on environment-wide evolution, not revolution.** Most organizations already have a complex set of initiatives, projects, and activities in progress for different parts of the business, so adding more could be daunting. The most valuable contribution of an environment-wide application strategy is to assess change in the context of the health of the whole organization. Evolution across the whole supports this goal more than revolutionary change in a few units. This approach does not mean outlier teams that revolutionize their processes, approach, or results do not have their place. In fact, those achievements should be celebrated as examples of improvement of the whole.
- ▶ **Think about systems and metrics just as much as goals.** Setting goals for improvements is often relatively straightforward. It is easy to say “we need to get better at X.” But goals of this nature, even with actions attached, may only attract a short-term burst of activity. A more effective approach is to select one or more metrics that capture the status of the environment in this dimension and work toward a set of behaviors and habits that both measures those metrics and moves the team forward.

Table 2. Sample metrics and processes for policy areas discussed in Part two

Recommendation	Reliability gains	Productivity gains
Decentralized authority	<ul style="list-style-type: none"> ▶ Decision time for changes or exception scored by the size of the issue 	<ul style="list-style-type: none"> ▶ Regular cadence of global and group planning processes to set policies ▶ Change request or exception processes between groups
Access everywhere	<ul style="list-style-type: none"> ▶ Availability of services and speed or service-level agreements (SLAs) per region, datacenter, or cloud 	<ul style="list-style-type: none"> ▶ Convention over configuration-tools and processes work the same way everywhere ▶ Transferable skills
Constraints = freedom	<ul style="list-style-type: none"> ▶ Adoption of design frameworks, patterns, or guidelines ▶ Compliance of solutions ▶ Number of exception requests 	<ul style="list-style-type: none"> ▶ Organization-wide sharing of interfaces, products, and internal services ▶ Regular cadence recognition of the adoption of central guidelines
Self-service	<ul style="list-style-type: none"> ▶ Speed or the number of steps to access key services or products with a metric such as "Time to first third-party user" for an API being used 	<ul style="list-style-type: none"> ▶ Regular cadence cycle of updated lists of available services and measure times required to access them by their users
Product mindset	<ul style="list-style-type: none"> ▶ The number of groups, services, and products using a particular product or service 	<ul style="list-style-type: none"> ▶ Regular cadence on the release of products and services, feedback mechanisms, means of communicating updates on changes, sharing usage metrics
Tend toward antifragility	<ul style="list-style-type: none"> ▶ Measure of time to heal after failures, as well as the (hopefully decreasing) trend on recurrence of similar failure 	<ul style="list-style-type: none"> ▶ Regular chaos tests inducing simulated or real failures across system to detect fragility ▶ Gradual widening of the envelope of the severity of system shocks induced
Automation everywhere	<ul style="list-style-type: none"> ▶ Measure the number of manual deployment or configurations or other steps—or total end-to-end time taken for full redeploys 	<ul style="list-style-type: none"> ▶ Regular review of automation metrics and ongoing program to reduce the largest pain points
Security in depth	<ul style="list-style-type: none"> ▶ Number of security incidents detected and caught ▶ Number of components adopting security guidelines and monitored 	<ul style="list-style-type: none"> ▶ Regular penetration and other security testing to stress system ▶ Post-action reviews

Part three: Architecture for success

The strategy principles in the previous sections are abstract so that they might be applied to a variety of systems. The next logical questions are:

- ▶ What do these systems look like?
- ▶ What are the components to consider, and how are they organized?

There is no single architecture that works for all cases. Architectural choices will always be unique to an organization, often representing core strengths to build upon. However, it is useful to provide a map of the typical elements of an application environment and how they connect when cloud-native and hybrid cloud technologies are applied.

In this section, we will provide an overall picture of what this means and drill down into a number of areas that have become critical to modern IT.

The big picture

Borrowing from our earlier definition of an application environment as the set of capabilities for application delivery and application development across the entire organization, a high-level view of a large, organization-wide application environment is illustrated in Figure 4.




Figure 4. High-level IT application environment

In this section, we focus on the components of an application environment and how they are connected. Each of these layers can be considered perpendicular to the strategy principles discussed earlier in Part two. The strategy principles should hold across these layers and inform their design. The division is as follows:

- 1. Platforms and application delivery.** The infrastructure that hosts the code processes that instantiate the applications that deliver value at the operating system and datacenter level.
- 2. Applications from custom code.** Support for the execution of custom code-based applications that are innovative in functionality and separate an organization from its competitors.

"Behind open source technology there's always a big community evolving the product very fast compared to proprietary technology."

Tobias Mohr

Head of Technology and Infrastructure, AVIATAR, Lufthansa Technik

8

4. **Applications from process automation.** Support for applications that involve not only executing code, but also processes that engage humans (workflows) and logical rules contributed by non-developer experts in the organization.
5. **Developer tooling, DevOps, and management.** Capabilities that wrap around the application environment to keep team members productive and systems functioning.

Each area is addressed in its own section to follow. In addition, addressing human processes and architectural patterns is a crucial success factor in implementing these strategies and architectures. Team dynamics, behaviors, and culture that maintain and evolve the system are all important factors for success. These processes and patterns also capture the organizational memory of the approaches that solve problems.

There are other ways to segment IT infrastructure depending on the focus, for example, mobile application support or Internet of Things (IoT). However, the model here is deliberately structured to illustrate the importance of the whole set of connected elements that make up an IT environment. In general, most IT strategy advice is either focused only on custom code development—ignoring that much of IT is about integration and processes—or assumes the adoption of a particular technology stack.

The application environments of most large organizations are vast and diverse, naturally spanning many forms of development. Typical organizations operate on the basis of thousands of individual applications, a wide range of SaaS services, commercial off-the-shelf solutions, and other systems. The architecture layers discussed provide a general overview of how to connect some of these systems. The examples refer to Red Hat technologies where relevant, but there are certainly other vendor solutions available. Any sufficiently complex system will need to rely on technology from many sources. As such, the specific products and technologies listed are for illustrative purposes only.

Platform and delivery

Improvements in platform and application delivery technology have arguably been the most obvious catalysts of cloud-native and hybrid cloud capabilities to date. The platform layers provide the substrate on which applications execute, integrate, and deliver value. The shift from the operating system to virtualization to containerized deployments now allows for a broad and deep spectrum of different kinds of deployment and automation. This shift in technology also furthers the idea that the exact same execution environment can be available to developers in an organization, no matter where they deploy, in the datacenter or on any one of a number of public or private clouds.

8 Red Hat case study, ["Lufthansa Technik builds hybrid cloud with Red Hat,"](#) May 2018.

In hybrid cloud environments with multiple on-premise locations, there is now an even greater premium on:

- ▶ Consistency of experience across locations, which ensures the same operating system versions, patch levels, and other system elements are available in as many of the environments as possible.
- ▶ Unified management across on-premise, virtualized, private, and public cloud systems, making it simpler to track the entire IT infrastructure in the same way.
- ▶ Control over security and compliance, which continues to be critical even as different cloud and datacenter deployments proliferate.
- ▶ Increased stability, agility, and availability. Often, the goal of infrastructure expansion is to add capacity for failover and redundancy. However, it is important to take a consistent approach to these deployments or outlier environments could hurt reliability.

In addition to creating a consistent platform across multiple cloud environments, the last few years have seen a new priority emerge: the use of containers and container platforms.

Containers have become the deployment infrastructure of choice in modern IT. These constructs package and isolate applications with their entire runtime environment—all of the files necessary to run. Rapid deployment and easy management via a container platform means that instances of applications can be deployed, updated, and scaled anywhere the container platform is available. Containers and container platforms truly bring the cloud-native computing vision to life on top of the hybrid cloud.

“Flexibility and speed are crucial factors in timely delivery of financial services to our customers. These significant outcomes of our digital transformation initiative were made possible by the framework we’ve built with Red Hat solutions.”

Chris Wang

Deputy General Manager,
Information Management
Office,
Cathay United Bank

9

In a hybrid cloud environment, using a container platform makes it possible to create a truly identical computer environment across all datacenter and cloud environments by:

- ▶ Providing identical compute infrastructure in each location.
- ▶ Allowing a packaged container to run anywhere within the hybrid cloud environment without modification.
- ▶ Helping operations teams see workloads across their various container clusters.
- ▶ Supporting the scalability of resources in each location.

The container platform creates an abstraction layer over the underlying resources, making them reusable and consistent everywhere. The Kubernetes-based Red Hat OpenShift® provides all of these features and more. The mix of locations in any customer’s hybrid cloud can now include hosted offerings on each of the three major public clouds—Amazon Web Services (AWS), Google Cloud Platform, and Microsoft Azure. More information on the use of containers for hybrid cloud can be found in the [Red Hat cloud containers e-book](#).

9 Red Hat case study. [“Cathay United Bank saves 50% of service development time with Red Hat,”](#) March 2019.

Three platform strategy considerations are:

- ▶ **Making the platform as self-service as possible** to allow different teams to access the capacity and services they need while having a tight and consistent security model. Allowing development teams to self-provision virtual resources is crucial for speed. However, self-provisioning needs backed-in controls to provide system security.
- ▶ **The reliability of the entire platform infrastructure across all locations.** This consideration involves looking not only at individual systems or datacenters but also at their interdependencies and change policies for when failures occur: Are systems simply "restored" or are they changed to prevent future failures?
- ▶ **Automation is key in complex hybrid cloud environments.** Developing automation is critical to keeping hybrid cloud management feasible. At the same time, as discussed in Part three, it is important to consider automation as code that must be managed. Otherwise, automation itself becomes a legacy problem.

Applications from custom code

Application code does not execute by itself. Executing source code is supported by substrates that have slowly evolved from early application servers to a mix of modern application servers, lightweight language runtimes, and container management platforms.




Figure 5. Execution of custom code in different contexts throughout the enterprise

As cloud-native technologies take hold, container infrastructure, in combination with language runtimes, provides the most flexible combination for innovation. As organizations adopt container technologies, this combination is likely to become the dominant deployment pattern with much of the previous application servers' orchestration capability being transferred to the container management layer.

For most organizations, this transition will take time, and it will likely have a mix of application hosting scenarios, including all of the approaches shown in Figure 5.

The following paradigms have emerged in application development and are being considered by the majority of large organizations:

- ▶ While Java and Javascript remain dominant in the enterprise for custom code implementations, the number of languages in use continues to grow.

- ▶ Tight integration and the inclusion of ancillary services for messaging and in-memory data grids are becoming important services in the environment so custom code projects can be efficient.
- ▶ Reactive programming is gaining significant traction as the need for rapid processing custom applications increases. Frameworks such as Quarkus and platforms like Apache Kafka have been seeing swift adoption.
- ▶ Function-as-a-Service (FaaS), first popularized through AWS's serverless Lambda approach and now having a more generalized compute approach, allows programmers not to worry about provisioning compute capacity. Instead, individual elements of code (functions) can simply be deployed and executed when the right events occur. FaaS solutions are now available on the major public clouds and can be provisioned (typically on a container platform) by on-premise operations teams for their own teams of developers.

These trends add significant new capabilities to enterprise IT for the development of custom code applications. Carefully selecting which technologies to adopt is important because the custom code systems, which implement the bulk of the differentiating functionality, make an organization uniquely competitive in the market.

"Our Red Hat operating environment frees us to focus far less on our infrastructure and spend more time innovating and delivering new products that improve our customer experience."

Michael Catuara

Director of Distributed Systems, TransUnion

10

Red Hat offers a broad application development portfolio that covers most of the areas already identified in this section. Some recent innovations include:

- ▶ A broad expansion of the enterprise-grade runtimes available to support custom code. Red Hat's all-inclusive runtimes bundle now includes Red Hat JBoss® Enterprise Application Platform, Node.js, MicroProfile, Quarkus, Spring Boot, and others. This bundle provides developers with the flexibility to choose the tools they prefer, while CIOs can continue to minimize risk across the whole spectrum of an organization's custom code.
- ▶ The inclusion of complementary services for messaging and in-memory data grid, with Red Hat AMQ and Red Hat Data Grid.
- ▶ Tight integration with Kubernetes and Red Hat OpenShift, delivering a clean separation of the lower-level deployment capabilities traditionally handled by an application server. These capabilities can now be handled in the container platform through the code-centric runtime support services in the runtimes packages.
- ▶ Support for Kubernetes-native function execution mechanisms that support FaaS wherever a Kubernetes container platform is deployed. Red Hat OpenShift (based on the [Knative project](#), which is a collaboration with Google, Red Hat, IBM and others) provides the substrate for messaging, eventing, and execution of a serverless approach to work across the hybrid cloud.
- ▶ With many organizations having both public and private cloud environments, they need application services not only in traditional on-premise datacenters but also in the cloud. Red Hat Cloud Services offers hosted and managed platform, application, and data services that help to reduce the operational cost and complexity of delivering cloud-native applications.

10 Red Hat success story. "[TransUnion modernizes IT, delivers a better customer experience](#)" November, 2021

- ▶ Lastly, in cloud-native environments, execution speed is critical. Applications that teams want to use in cloud, container, or even serverless environments should start up and shut down in a fraction of a second: scaling to zero when they are not needed while being available nearly instantly when called upon. The Red Hat build of [Quarkus](#) now provides this speed and agility for Java(TM) Virtual Machine (JVM) languages.

Key considerations of application runtime strategy are:

- ▶ **Autonomy with consistency**—supporting the use of different languages and approaches throughout the organization.
- ▶ **Tight integration with DevOps and developer tools**, which is important for developer productivity. Automation of standard processes, checks, and balances allow more code to be written, tested, and deployed efficiently.
- ▶ **Establishing a comprehensive security strategy**, which ensures that, from the operating system through virtualization and from the container layer through to the code runtime, all systems are comprehensively updated by automated rollouts from trusted vendors.

"We picked OpenShift because it is highly scalable, it's an enterprise-ready container platform with powerful orchestration Kubernetes. It has become a foundational element for us."¹¹

Ken Finnerty
President of Information
Technology, UPS

Integration is critical

Since IT systems were first connected to networks and operating systems first allowed interprocess communication, integration capabilities have been necessary. The number of ways software systems can be integrated has ballooned. Integration has become one of the most critical parts of the IT stack.

Since application environments now span multiple datacenters and clouds, integration solutions also need to adapt to modern requirements. While the traditional patterns for integration, such as enterprise service bus (ESB) deployments, have brought organizations a long way, they are far from sufficient for hybrid cloud challenges. In particular, there are a number of challenges integration technology needs to meet:

- ▶ Uninterrupted operation across multiple locations. This operation means that messaging, API management, transformation, data mapping, and other capabilities should all be available in multiple physical locations. The capabilities also need to work together across locations. For example, messaging solutions often need to be used across multiple physical locations, not just within a location.
- ▶ Integration as code. Integrations between systems perform critical transformation and data mapping in real time. This capability requires code changes to the software endpoints, delivering data that can imply synchronized changes to the integration. Integrations should be considered code, just like applications, so that they may be versioned in the same way.
- ▶ Extreme scale. While integration technology started in the datacenter connecting primarily back-office applications, in many cases, customer-facing applications need access to those same back-office systems. Meanwhile, traffic volumes have increased significantly. Integrations, therefore, need to scale rapidly and cost-effectively for the most heavily used data and transaction systems.

¹¹ Red Hat Summit keynote recording. [UPS keynote](#), 9 May 2019.

- ▶ Serving more stakeholders. The number of systems has grown rapidly, and IT has infiltrated almost every business function. The number and type of integrations required in an organization have dramatically increased. Supporting integration activities is becoming increasingly critical not only for development and operations teams, but also less technically trained individuals in the line of business areas, known as citizen integrators.
- ▶ The increasing importance of data. In line with their use in customer-facing applications, integrations are also front and center in the trend toward large-scale, real-time data capture and analysis. Moreover, the proliferation of IoT devices requires new ways to manage the vast amount of information they generate. Data needs to be processed, stored, and often replicated accurately for modern systems to operate effectively.

To cope with these changes, integration technologies themselves are evolving. As they combine with cloud-native capabilities, integration technologies in the Red Hat portfolio:

- ▶ Can be deployed in a fully distributed manner, with individual integrations wherever they are needed.
- ▶ Are container-native, which permits automated capacity scaling.
- ▶ Can be used on demand with instant startup and shut down via container-based Knative technology.
- ▶ Are manipulatable as code, which means normal DevOps practices can be applied and are accessible via a graphical drag-and-drop interface. This capability allows highly skilled and less-skilled individuals to collaborate on the same integrations.
- ▶ Include an even broader range of messaging capabilities, from traditional AMQ-style messaging to Apache Kafka.
- ▶ Include messaging, API, and integration technologies that can be deployed across different clouds and provide the application-level communication needed for end-user applications to operate.

The most important integration considerations for organizations are:

- ▶ Mapping out a strategy for which products and services become available and how they will be accessed by other software teams. Often APIs are an important part of this broader adoption.
- ▶ Determining a strategy for who owns integrations between these systems. Distributing this workload increases the autonomy of individual groups and reduces the workload on centralized Integration Competency Center (ICC) teams.
- ▶ Addressing strategic questions, such as setting best practices for changed-data capture, virtualization of federated data, messaging, and data transformation across datacenter and cloud integrations. How will these capabilities connect with the platform layer and provide seamless service everywhere?

A deeper dive into integration strategy in cloud-native and hybrid cloud environments can be found in Red Hat's [Agile integration e-book](#).

Applications from process automation

Software touches nearly every business process an organization executes. Many of these processes involve people (employees, partners, and customers) and also complex decisions based on data and states.

The third type of application that is radically changing in the hybrid cloud era is process-driven applications.

There is a clear, growing need for applications that can automate entire business processes, and make intelligent business decisions at appropriate points along the way.

As IT spreads infrastructure through the hybrid cloud and influence throughout their organization, the following trends are emerging for process-driven applications:

- ▶ **The adoption of microservices patterns, APIs, and integrations** that creates an opportunity to connect processes with new systems throughout the enterprise—and build richer workflows and apply automation to data sets or systems that were not previously accessible.
- ▶ **Solutions using configuration, business-level logic**, and rules are becoming increasingly important. Although custom applications deliver much of the novel functionality available in the environment, the ability to call that custom logic in ways that reflect changing business needs is essential. Configurable rules are capable of executing actions as FaaS invocations in the cloud and containers.
- ▶ **An increase in process distribution**, as the number of datacenter and cloud locations in an organization grows, is needed to span these locations, requiring a cloud-native approach to process management. This approach is not simply deployed in one location, but it can be configured to run wherever it is needed and create consistent flows across all locations.
- ▶ **The rise of BizDevOps** that aims to combine business needs directly with agile application development.
- ▶ **Scalability.** Process-driven applications, particularly in organizations with large numbers of employees or customers, now underpin some of the most highly used IT services. As a result, it has become necessary to scale up and down usage of key processes on demand to avoid failures and bottlenecks.
- ▶ **The use of serverless and FaaS platforms** by business rules and business process management systems that helps with scalability, performance, and cost savings. Decisions and their actions can be executed as FaaS. Also, steps in business processes can be executed as FaaS.
- ▶ **The confluence of IT integration** between back-office systems and customer-facing systems (and the resulting growth in data) has created an enormous opportunity for organizations to begin deploying intelligent applications for their clients. Making rule-based decisions, splicing together human and automated decision-making, and doing this with speed and scale are all key to taking advantage of these opportunities.

"We want to go on this digital transformation journey with a partner that can help us create the platform that we need—sharing code, information, and ideas. Using open source implies the use of common standards that guarantees platform stability."

José María Ruesta

Global Head of Infrastructure, Service, and Open Systems, BBVA

12

12 Red Hat success story. "[BBVA builds fully automated cloud platform to better support customers](#)," May 2018.

"The direction of Red Hat Ansible Automation Platform enables us to execute automation in containerized environments, supporting new strategies for innovation in the business. Now, we can scale automation systematically across the enterprise with consistency and as demand requires to hit those two priorities with fewer barriers, using a platform that's already integrated into our workflow."

Michael McCarthy

Delivery Architect, Production Engineering Group, Gamesys¹³

Each of these trends means that a modern wave of process-driven applications is now a crucial component of IT strategy.

As with integration technologies, cloud-native development is also making major changes in process automation technologies possible. In the Red Hat portfolio, it means:

- ▶ **Native deployment on containers.** Tight integration with containers and container orchestration technology is a core part of Red Hat's process automation platforms. Process-driven applications benefit from the scalability, resilience, and location independence provided by a cloud-native architecture.
- ▶ **Serverless and FaaS.** Red Hat's process automation solution, based on technology from the open source [Kogito](#) project, extends cloud-native capabilities by supporting serverless and FaaS models for process-driven applications. Kogito is built from the ground up with Quarkus to provide the rapid startup, scaling, and memory efficiency needed for on-demand services.
- ▶ **Cognitive computing.** Cloud-native architecture provides the scale necessary to support resource-intensive artificial intelligence (AI) and machine learning (ML) workloads, which can extend decision services to cover a wider range of scenarios. By incorporating predictive models with user-written Decision Model and Notation (DMN) logic, Red Hat Decision Manager helps organizations increase automation and reduce exceptions requiring human input.
- ▶ **Intelligent management.** Cloud-native development reduces operations workloads through Kubernetes Operators, which automate application life cycle management and fault response. Red Hat process automation solutions include Operators for process-driven applications to easily manage distributed hybrid cloud workloads.

The strategic decisions relevant to this area include:

- ▶ **Allowing more employees, partners, and customers to participate in the creation and configuration of process-driven applications.** This type of application creates an opportunity for decentralized authority where technical specialists can create a safe, structured environment. Domain experts can then use this to control the applications for their own needs.
- ▶ **Determining protocols and procedures, roles and responsibilities** for collaboration between line-of-business and IT when automating manual business processes.
- ▶ **Deciding on approaches to use FaaS and cloud computing for decision and intelligent process automation (IPA).** For example, whether to use IPA-as-a-Service in a cloud environment, on-premise, or in a combination of both.
- ▶ **Providing process management on-demand as a service, which may bring the greatest reward** by allowing teams that previously did not have process specialists to begin capturing and instrumenting their own workflows.
- ▶ **Moving toward antifragility.** By replacing custom-engineered, compiled code for processes with configurable rules and standard formats, complex processes can be better validated, and change does not need to involve code deployment. This approach reduces the risk of system breakage and helps more team members understand the dependencies inherent within important processes.

¹³ Red Hat press release. "[Red Hat Ansible Automation Platform 2 Drives Cloud-Native Automation and Helps Developers Become Automators](#)," 29 September 2021.

"Red Hat Decision Manager plays an important role in Northwell's digital business strategy and has enabled us to accelerate the development of core healthcare applications."

Vipul Kashyap

Director of Clinical Information Systems and Enterprise Information Architect, Northwell Health

Developer tooling, DevOps, and application management

DevOps and developer tooling represent the foundation of cloud-native technology. In essence, many cloud-native technologies have evolved from the core need for infrastructure to support developers more efficiently and effectively. The Cloud Native Computing Foundation's (CNCF) definition of cloud-native technology starts with:¹⁴

- ▶ Cloud native technologies empower organizations to build and run scalable applications in modern, dynamic environments such as public, private, and hybrid clouds.

CNCF covers some examples and then highlights:

- ▶ These techniques enable loosely coupled systems that are resilient, manageable, and observable. Combined with robust automation, they allow engineers to make high-impact changes frequently and predictably with minimal toil.

This combination of resilience, manageability, and observability, (from the operations side) and engineering impact (from the developer and engineering side) has fostered a number of technologies that have come together under the cloud-native banner.

These attributes are what makes cloud-native technology a compelling technology set used to address many of the challenges that emerge as organizations move along their trajectory of adopting multiple cloud environments and datacenters as their infrastructure footprint.

DevOps practices, application services, developer tools, and management, as well as new cloud-native variants, all comprise and support the IT environment of an organization. They work together to create the organization's application environment. In turn, that application environment supports the functionality that allows employees, partners, and customers to derive the value they need.

Much has been written about the successful application of DevOps techniques in organizations.

[Red Hat Open Innovation Labs](#) and other services teams provide programs to help get started.

Rather than repeat this general guidance, this section focuses on various trends accelerated by hybrid cloud adoption:

- ▶ The notion of infrastructure as code, which comes from cloud-native DevOps, is becoming critical as infrastructure and applications span multiple datacenters and cloud environments. Without the ability to fully automate deployments in multiple locations, hybrid cloud environments become increasingly hard to manage.
- ▶ This need for automation extends to application code as well. Since applications are now tightly woven integrations across numerous services and APIs in multiple locations, deployments need to be carefully orchestrated to avoid downtime.

¹⁴ From the Cloud Native Computing Foundation (CNCF) official definition of cloud native v1.0 <https://github.com/cncf/toc/blob/master/DEFINITION.md>

- ▶ Almost all IT environments are mixed platform environments that can include bare-metal deployed machines running a variety of operating systems, virtualized server real estate, private cloud, public cloud, and container platforms. The mixture of these different platform levels means that any holistic approach needs to connect management practices across platforms.
- ▶ Development tools need to change and update with infrastructure, providing live testing and other cloud-native features. Cloud-integrated development environments, in which code lives entirely in the browser, are gaining traction to meet this need.
- ▶ Observability and tracing becomes harder and more crucial in hybrid cloud environments. Code execution often touches multiple locations and must be tracked across each invocation. Failures in disparate parts of the system need to be correlated to determine what went wrong.

Each of these trends raises new challenges, but today's tools are evolving rapidly to cope with these struggles. These challenges all result from an increasing decentralization of infrastructure across locations, application types, and technologies. Part of the response is to standardize elements of deployments. A second part is to deploy tools and architecture patterns that explicitly solve for the distributed nature of the system. Both are needed for success.

Some of the developments in Red Hat's solution stack include:

- ▶ Multicloud management in Red Hat OpenShift, as well as a variety of tools for integrating continuous integration and continuous delivery (CI/CD) pipelines with container deployments that create a powerful hybrid cloud environment.
- ▶ [Red Hat CodeReady Workspaces](#) provides a completely web browser-based integrated development environment (IDE) optimized for container-deployed code. Features include container-based workspaces, prebuilt or custom stacks, centralized management of developer configurations, and much more.
- ▶ Since many users have both public and private clouds, they need application services in traditional on-premise datacenters and in the cloud. [Red Hat Cloud Services](#) includes a family of application services that reduce the operational cost and complexity of delivering cloud-native applications, and allow development teams to focus on core competencies.
- ▶ Red Hat Ansible® Automation Platform 2 works in a world where hybrid and multicloud computing have become critical IT components. This solution gives IT organizations an automation platform that can bridge traditional systems with these modern services, stretching from the enterprise datacenter to the remote edges of the network.¹⁵
- ▶ Red Hat's predictive monitoring solution, [Red Hat Insights](#), is now included with all Red Hat Enterprise Linux® subscriptions. The platform provides integrated security, performance, availability, and stability across a variety of Red Hat platforms to increase visibility in the hybrid cloud.

¹⁵ Red Hat press release. "[Red Hat Ansible Automation Platform 2 Drives Cloud-Native Automation and Helps Developers Become Automators](#)," 29 September 2021, .

From a long-term planning perspective, the following strategic areas are key for developer tooling, DevOps, and application management:

- ▶ **Universal capabilities** ensure that the same underlying infrastructure, services, and automation are available in each location and for the application delivery infrastructure (platform and automation) as well as the applications themselves. Often deployments need to run the latest version of each app in multiple locations.
- ▶ **Developer tooling approach.** For example, a browser-based IDE running on containers may be the best option for providing a consistent developer experience across on-premise and cloud environments.
- ▶ **A product mindset among development and operations teams.** Creating applications and services for reuse is hard work. However, it is harder to build trust in products and services operated by others. A product mindset requires teams to think long term about the systems they provide. This approach impacts their planning process and will need to be captured in the automation, daily habits, and practices that go with updating, deploying, and retiring code.
- ▶ **Automation.** Without careful management of the processes and automation code themselves, a successful DevOps implementation quickly becomes a rigid system with diminishing returns.

Part four: The application environment journey

The first thing to understand about changing the application environment is that it is hard. Aligning data across multiple sites and public cloud products is not a trivial undertaking, and sometimes after the process is completed, results can be suboptimal, which means the process must be tried again.

Organizations need to approach the process as a journey and not a one-off project because sometimes the results are not linear.

Except for newly established startups, there are very few situations where IT strategy is purely about the new. Typically, IT strategy is a study in how to evolve the entire functioning system for improved reliability and productivity. Individual bursts of fast productivity may sometimes work, but without the general use of best practices, these instances end up draining resources.

This section covers the journey of cloud-native and hybrid cloud application environments by sequencing the strategy areas from Part three into Figure 6.




Figure 6. Linear representation of the evolution of resilience and productivity in application environments

Note that these are not technology adoption stages. There may be technologies that are never adopted in some (or many) parts of a particular organization. The stages reflect policies, capabilities, and decision-making. In fact, there are some technologies or solutions that might be adopted that could need to be reworked or re-implemented, and the time to assess and circle back on these changes should be factored into the process of adoption.

Expecting large, complex IT organizations to advance all the systems, processes, and strategic directions at once or in lockstep across the organization is also unrealistic, and attempting to do so could even be damaging to the solution you are trying to implement.

Visualizing these elements as a sequence is designed to highlight the logical coordination of improvements and, more importantly, emphasize the whole application environment as a single entity. Improvements can be made in any unit, at any stage, but it is useful to consider how those changes affect the whole environment: Will they spread? Or will they stay segmented and disappear?

Fragmented bureaucracy

Fragmented bureaucracy might seem like an unkind description of an organization. However, given the complexity of today's IT systems, this is how many IT organizations operate. The rules, processes, and technology choices of previous builder generations are what made the systems and the organization. These decisions are reflected in the IT organization. It is also unlikely that any set of changes, no matter how great or how persistent, will make the organization and its systems perfect. Any moment of perfection would become obsolete by the next change. Technical debt is seen as a top barrier preventing success in digital transformation for upwards of 23% of companies.¹⁶

In fact, IT organizations are not just implementing an application—they are always implementing a complex ecosystem with a potential multidecade life span.

It is important to not just get away from the label, but to move the organization and structures forward to improve. Bureaucracies are required structures to operate such large, complex organizations of people and technical systems—some fragmentation will always be present and even necessary. The questions to consider are:

- ▶ How good is the current structure and the application environment that it creates for the organization?
- ▶ What are the most important directions and mechanisms of change to keep the whole improving?

Steps along the cloud-native and hybrid cloud path

For many IT organizations, using multiple datacenters and multiple public cloud locations, often from more than one cloud provider, is now a daily reality. From a technology perspective, there are a variety of strategies to pursue that help make such environments more manageable.

¹⁶ Red Hat overview. "[2022 Global Tech Outlook](#)," 16 November 2021.

Improving the application environment that powers your organization might require a range of technology, but above all, it requires making strategic decisions about prioritizing areas of change. When viewed from that perspective, the following four areas look like steps toward prioritizing and implementing long-term change:

- ▶ **Product mindset.** Adopting a product mindset and cultivating it for any system that is intended for reuse helps teams act more autonomously. Having this mindset is also a prerequisite for teams in the organization to rely on and trust systems run by others. In fact, this focus is also a characteristic of a microservice architecture that advocates that a team should own a product over its lifetime, from development to production maintenance, resulting in a closer linkage to business capabilities and how they can be enhanced. A product mindset helps guide the discourse on which interfaces and processes are truly important. Where are the domain boundaries and which APIs and functionality sets need solid commitments behind them?
- ▶ **Services everywhere.** Emphasize the need to think beyond initial deployments and plan for uninterrupted functionality to be available in multiple locations. There can be costs to this planning (or rollout) but they are usually lower than dealing with wildly different solutions to the same problem. This area also encompasses the push for self-service where possible so developers, and potentially users, can self-provision resources. There are significant efficiency and productivity gains in running an infrastructure in this way. These gains will grow in importance as container and serverless technologies become more prevalent.
- ▶ **Automation as code.** Raising the level of automation across the IT environment provides a higher return on investment, but only if this automation is explicitly maintained in the same way code and applications are maintained. Without explicit management, today's automation may become tomorrow's awkward legacy. A similar mindset shift applies to integrations, decision-making rules, and process automation code. These generate or underpin application functionality. In many cases, more integration code is present in a system than custom logic. Explicitly considering these systems as code makes it simpler to manage deployments across systems.
- ▶ **Antifragility.** The move toward antifragility represents the presence of strong processes (automated or human) that react to failures and improve the system beyond where it was before. Few organizations have adopted extreme, Netflix-like chaos engineering at any significant scale, so progress towards antifragility is a valuable goal. Similarly, including security in this and all previous steps represents a significant leap toward proactivity.

Is there a methodology?

There is no formal methodology to accompany this e-book because the topic of large-scale IT is too complex to discuss in this manner. Rather, the objective is to draw together several key strategic themes about how hybrid cloud and cloud-native technologies are impacting organizations.

While keeping these strategy questions in mind, thinking of an organization's IT systems as a single application environment brings into focus a number of considerations for the reader:

- ▶ Which strategic dimensions will deliver the most gain?
- ▶ How do we spread good practice in one area or across the organization?

- ▶ How do key types of platforms (operating systems, virtualization, containers) and applications (custom code, integrations, process-driven apps) fit together?
- ▶ How do all of these concepts relate to developer productivity and stability for operations?

For a richer analysis of these topics, or an assessment of the areas of greatest impact, please see the next section for programs run by the Red Hat Services teams.

Part five: Conclusions and next steps

IT systems have never been so valuable. However, without a cloud-native, agile strategy that addresses both reliability and productivity, businesses will increasingly struggle to compete in the market. While previous eras of IT have experienced great innovations, the past few years have seen an unprecedented reliance on IT to power almost every business function. Back-office and front-office systems must work together like never before.

The desire to explore new technologies and operating models needs to be weighed against the fact that development teams are stretched thin and often overwhelmed just maintaining applications.

To succeed in this environment, organizations need to provide a healthy, robust, and productive application environment that serves as a foundation for innovation. This application environment needs to span multiple datacenters and clouds while weaving in the latest, most productive cloud-native technologies.

Your organization needs applications that distinguish you from competitors while engaging stakeholders and providing value. Our objective is to consolidate strategic insights from customers and technology deployments to help others plan and strategize, and to offer services that help with development, delivery, integration, and automation.

While not all the insights in this e-book will apply to every scenario, hopefully a number of the questions and considerations posed will be helpful on your IT journey.

Discover more about how Red Hat can support you in your cloud-native and hybrid cloud application journeys:

- ▶ See how Red Hat Consulting can help—get best practice and planning guidance with a [Consulting discovery session](#).
- ▶ Check out our [Services Speak blog](#) for insights, tips, and more.

Learn more information about Red Hat's technology offerings:

- ▶ Application delivery: [Red Hat OpenShift](#)
- ▶ Application development: [Red Hat Application Services](#)
- ▶ Developer tools: [Red Hat Developer Toolset](#)
- ▶ Management: [Red Hat Smart Management](#)



About Red Hat

Red Hat is the world's leading provider of enterprise open source software solutions, using a community-powered approach to deliver reliable and high-performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers integrate new and existing IT applications, develop cloud-native applications, standardize on our industry-leading operating system, and automate, secure, and manage complex environments. [Award-winning](#) support, training, and consulting services make Red Hat a [trusted adviser to the Fortune 500](#). As a strategic partner to cloud providers, system integrators, application vendors, customers, and open source communities, Red Hat can help organizations prepare for the digital future.

North America

1888 REDHAT1
www.redhat.com

Europe, Middle East, and Africa

00800 7334 2835
europe@redhat.com

Asia Pacific

+65 6490 4200
apac@redhat.com

Latin America

+54 11 4329 7300
info-latam@redhat.com

 facebook.com/redhatinc

 [@RedHat](https://twitter.com/RedHat)

 linkedin.com/company/red-hat

redhat.com
#000000_0021

Copyright © 2022 Red Hat, Inc. Red Hat, the Red Hat logo, OpenShift, JBoss, and Ansible are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. The OpenStack word mark and the Square O Design, together or apart, are trademarks or registered trademarks of OpenStack Foundation in the United States and other countries, and are used with the OpenStack Foundation's permission. Red Hat, Inc. is not affiliated with, endorsed by, or sponsored by the OpenStack Foundation or the OpenStack community. Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle America, Inc. in the U.S. and other countries. All other trademarks are the property of their respective owners.